



ISRCS Panel Discussion Notes

August 13, 2009

1. Topics:
 - a. Human systems and resilience, what are the issues and how do we address them
 - b. Cyber awareness and resilience, what are the issues and how do we address them
2. Panel Members:
 - a. Moderator: Margie Jeffs, Idaho National Laboratory
 - b. Session Chairs: Miles McQueen, cyber researcher, Idaho National Laboratory and David Gertman, human systems researcher, Idaho National Laboratory
 - c. Annarita Giani, postdoc, University of California Berkeley
 - d. Eugene Santos, professor, Dartmouth College
 - e. Ron Boring, researcher, Sandia National Laboratory
 - f. Johanna Oxstrand, researcher, Vattenfall Ringhals AB
 - g. Bob Richards, researcher, Idaho National Laboratory
 - h. David Woods, professor, Ohio State University
3. Approach
 - a. Overview of Format (6 Questions will be asked)
 - b. Introduction of panel members by Session Chairs
 - c. Recitation of question, then panel and audience response
 - i. Presentation of Question and Example by Moderator
 - ii. Response to Question by Panelists
 - iii. Audience Response to Question
 - d. Panelist Participation
 - i. Moderator can interrupt panelist if individual is going over the time allotted, given a 1 minute grace period
 - e. Audience Participation
 - i. Moderator picks individuals asking questions to fill 15 minute period
 - ii. Audience participation at this point is only to address question, not ask further questions
 - iii. Moderator can interrupt and end response for individuals asking questions, usually under conditions of time, providing off-point discussion, or being redundant.
 - f. Final Discussion
 - i. Panelist given a few minutes to provide closing comments
 - ii. Audience given 20 minutes to ask questions of the panelists

Panel Discussion Questions:

1. Do you think human system concepts or models can be used to cover system resilience? (Human Systems)

Panelists

- Yes, absolutely. Concepts in human systems are important. Need to think of as a system, not independently.
- Stop thinking of as people OR machines OR automation OR.... This (resilience engineering) is a paradigm shift to see these things as a system. We should go back to the control models that John Doyle talks about. Synchronization has to be in the forefront of our thinking. The human is the only competence model. As we scale up systems, focus of attention should become a focus area.
- Human system concepts are necessary, but not sufficient. We don't have a (human systems) model that takes everything into account and are therefore surprised by what occurs in operating events, cyber or otherwise. Additional factors are needed to get a complete picture of resilience.
- Models can help to start addressing resilience. The strength of the human is to react quickly to a surprise situation. The vision for the models/concepts would be to capture that human ability to adapt and react to surprise.
- Not sure how human system models translate in this surprise arena.
- As we talk about the human side, as we go up to social systems, it is difficult for engineering systems to have a good communication system among themselves let alone with management. How can we grow this into meta systems?
- We could start thinking about social systems – maybe layered, maybe not. Academia is not. Academia could do a better job about thinking about how to get out of stovepipes. What are elements of communication? We should go back and look at case studies (and there are many good case studies from which to select) to decide how to do this better. There are many good case studies from which to select (e.g., from doctors in health care system who have written about how to do it right).
- We should be able to offer guidance in this area.

2. Can resilience bridge the gap between hardware reliability, software reliability, cyber awareness and human reliability? (Human Systems)

Panelists

- If resilience could come to the party and be the glue, that would be a great thing but that seems to be a tall order. The resilience framework may be a different way of looking at things, a different approach, and could potentially be useful but it will take significant development. Hardware reliability and human reliability (HRA) are joined probabilistic risk analysis), there is only a little hand waving over software reliability. Cyber events are ignored in these studies. The estimates are in probabilities, it is hard to know whether resilience would be a measure included in the same report or modify the reliabilities. It could be another overall systems measure.
- For now, would rather keep resilience as a separate concept. Reliability has its place. We shouldn't even think of resilience as bridging the gap between these things.

3. Can you think of a compelling example of either system resilience or brittleness? (Human Systems)

This question was not addressed due to time constraints, except in closing remarks.

4. Can the malicious actor be effectively modeled, or is it best to confound and randomize? (Cyber awareness)

Panelists

- Yes, the malicious actor can be modeled. Some modeling is needed because otherwise it is not helpful to talk about solutions. In security in general, the vendors don't tell you what threat they are mitigating with any specificity. The wrong answer is the "find and fix vulnerabilities" model. The right answer may be to put the attacker in a difficult position. Right now, attackers work in the weeds looking for one weakness so we need to use complexity against the attackers.
- For example, if we postulate an attack/attacker is attempting to blow up the top 30 refineries, it would be good to have random operating times for certain functions in those facilities. Something like this is feasible and practical. The attacker needs to be pushed into the probabilistic world.
- The answer to the question is no and no. Model the co-adaptive process instead. If we are not careful, we may enhance the attacker's capabilities or create vulnerabilities in other areas. There needs to be a co-adaptive changing process. Then, we will be able to sustain the ability to carry out defensive strategy. We also need to understand how to mingle and mix cooperative and competitiveness. When we model our adversaries, we usually misread their intent and end up with a dynamic management problem.
- Does not think his points disagree with David's points. It is a co-adaptive system. Agrees there is opportunity for making things far more difficult.
- The structure of model can be designed. However, designing very precise models is very hard.
- The malicious actor can probably be effectively modeled if the actor is considered as part of a co-adaptive system. The older risk model approaches, that present a few physical barriers to be defeated by an attacker, are limited. Particularly if we are speaking about cyber systems. There is a huge challenge in getting the human systems model right because of uncertainty regarding the extent and timing of adversary behaviors. There is a sense you should do something in the interim. This effort to make things difficult for the attacker could include randomizing and confounding methods.
- Trying to model the attacker turns HRA on its head. We are trying to design a system with what could get in the way of an attacker attempting to get into the system. How can we increase their likelihood of errors? Would be modeling to confound.
- Another way of looking at this would be to define various attributes and look the effect. Then prune down.
- In reality, we are so far from effectively modeling the malicious actor. We first need to look at "what is the system vulnerable to?" and then, from our perspectives ask "have we closed all the obvious stupid vulnerabilities?" We need to work on these immediately – we have so much opportunity in this area. This may be a naïve, abstract model of an adversary but we need to start somewhere.
- We need to start considering a more adversarial model and begin understanding their intent.
- As our models get better, does it necessarily mean our resilience is getting better? Modeling is

important, but what do we do with it once we have it? How do you increase adaptive and capacity factors? Do we know how to effectively do that for complex situations?

- We can't just chase the areas where we have holes all the time. We don't have infinite resources or infinite time. We need to take a more active strategy requiring a dramatic cognitive shift. Can't simplify and make less dynamic – this is not good enough. Must be adaptive and flexible.

Audience

- Start with a model of the attacker's strengths. If we want to focus on targeted attacks, we need to focus on what they may do to you.

5. Can system attributes be effectively randomized to take away the advantage of the malicious attacker? (Cyber awareness)

Panelists

- One of the things related to this is fixed points in a network. So if we make it easy to plug in application and hardware, we make it easy for attack. We need to minimize constraints. Hacks into systems give us a clue as to what the architecture should be. So the first step should be to get constraints down to a small number and then put randomization around them.
- We are being hindered more and more by not being able to find the relevant stuff. We have compartmentalized our systems to prevent security breaches and then in turn tried to connect these. We need to think about interactions across different levels.

Audience

- Randomization is just one aspect. Whenever you randomize, you must always have one thing that is fixed. This does not necessarily make the attacker's job more difficult but it may make programming more difficult.
- From his experience, there aren't too many cases that can't be overcome. Our goal is to make their lives more difficult.

6. How should Human-in-the-Loop (HITL) be accounted for in cyber security of control systems? (Cyber awareness)

Panelists

- As we expand the definition of attack (cyber and physical), it is hard to avoid modeling human in the loop for both sides (attackers and defenders). Not sure how to account for HITL but we must do so at some point. We take care of this in safety studies but not sure how to take care of this in resilience. The fault tree approach may not provide us with the right answers.
- We have tried to account for the human in security situation by instituting things like a 2-person rule and looking at cyber potential via IDS and other things. However, we are missing the big picture of the insider threat because we are so worried about following confining rules and regulations and are completely missing the bigger challenge of some insider with access to cyber or physical systems.
- We keep layering things on (procedures) for honest people who wouldn't corrupt the system anyway. How do we really deal with those who are dishonest and are an insider threat?
- The burden is on us to understand techniques to apply when deal with generic tradeoffs. You can play games with rules and will be migrating randomly in space. We need some opportunities to study these real life cases.

- Cited an example within a hospital emergency room. We have bad polycentric architecture. We need to change that.
- Any action that a human takes...the adversary will incorporate that into their attack plan. There needs to be a long-term drive into better architectures. Need better awareness of when and when not to pull the plug.

Audience

- Observability and categorization are important. Demonstrating intent is an important component for a system owner to understand so they know when to NOT operate the system. Where do we draw the lines and make these decisions? How does a human being get the right input to make these decisions so they know when to not operate the system (which sometimes seems counterintuitive)?

Key Messages/Takeaways

Panelists

- Examples of resilience or brittleness – we need to develop a portfolio of examples for each. There is a good start in book “Resilience Engineering.” Thought about Valdez example – there was a setup for failure. Production was emphasized, better, faster, cheaper mentality causes us to miss the bigger global safety picture. A two person crew, working double shifts, one very junior, one with a substance abuse problem, at night, taking risks, poor communications with the coast guards was a recipe for disaster. The system was brittle beyond belief. Traveling near Bligh Reef had become a local optimization strategy at a global expense. We need to attack problems and questioning from a process and resilience perspective. Tokai Mura was another example where there had been an erosion of knowledge and procedures and a production mentality that allowed for bypassing criticality safe geometry. It ended with criticality from double or triple batching using a steel bucket. It is beyond belief. People doing well and stretching the system for productivity gains were setting system up for failure.
- The main reason for attending symposium was to learn more about resilience and figure out why it is not the same as safety culture but getting more confused. They are trying to incorporate all these things into safety culture. Thought safety culture equaled resilience. Resilience seems too broad for us to handle right now and can’t get our heads around it.
- The reason for participating in this symposium was to clarify resilience concepts. Is still struggling to reconcile Human Reliability and resilience. Would have liked to address Question #3 (“Can you think of a compelling example of either system resilience or brittleness?”). Most of HR is for NRC. I have had the opportunity to run simulation of steam generator rupture with American and Swedish crews. Had HR people running through the analysis – Swedish crews had stopped work and had a discussion about how to proceed. This surprised the Americans because they broke away from procedures. Procedures are considered “the law” but this example showed that in a complex incident, it was better to break away from the procedures. The humans were adapting to the situation and were more resilient.
- Good to see the diversity of this group. As each of us approaches the problem, we need to remember to keep the whole spectrum/viewpoints in mind. We need to keep our views broad so don’t get lost in our own, potentially narrow, perspectives.
- Don’t believe we are using humans properly/fully in systems. A TV program illustrated resiliency in humans. In the program, top athletes were encumbered in various manners (e.g., a basketball player was blindfolded during free throws). Though blindfolded, the athlete still

sank 8 of 10 free throws.

- Played simple video filmed in Bangkok which demonstrates adaptivity of humans. (Train tracks were quickly converted into a market after train has moved through the area). When try to explain people, remember they are adaptive agents and then build on that.

Q&A

Audience

- Question: What about economic incentives? They drive us to be less adaptive. Do we need public policy, etc.?
- Panelist Answer: Regarding incentives in general, accountability systems are negative. For example, the prescribed fire discussed earlier in the symposium was a great example of people looking in hindsight and then blaming. Very worried about how tend to fall back into these negative ways of behaving. Yes, it would be helpful to have public policy.